Fundamentals of Power Electronics: A Narrative Illustration

By Will Mefford

What is Power Electronics?

Power Electronics is a specialized field within the umbrella of Electrical Engineering that seeks to efficiently control and/or convert power using switches and passive storage elements. The oldest example of a power electronic system is the diode rectifier, which converts an AC voltage source into an unregulated DC voltage source using only naturally commutated switches. As devices more sophisticated than diodes became available, more topologies and control schemes became possible, allowing engineers to convert from one source to another in terms of voltage, current, and frequency all with efficiencies on the order of 90%. This flexibility permits the efficient and precise control of motors, power grids, lighting systems, heaters, and more.

The Problem

Let's say as an example, a young budding engineer would like to move a small trailer loaded with gravel from one side of his yard to the other, but he is particularly work averse and has only his daughter's broken toy car to provide a tractive force. If only he could repair it well enough, he could spare himself the misery of breaking a sweat, or worse, trying to repair the lawn mower. The electric motor and battery both seem in fine condition, but the magic box under the hood seems to have become a bit crispy after being left in the last torrential downpour. Connecting the motor directly to the battery might work, but this doesn't really afford much control over the speed, and since his progeny will be behind the wheel, some control would be preferred over none at all. At first glance this seems simple enough—DC motors are well behaved power converters with friendly proportional relationships between voltage, current, speed, and torque. But how will the battery's voltage be adjusted to in turn adjust the speed? If the motor is small enough and the engineer green enough, he might be tempted to use a particularly beefy opamp. Or perhaps a resistor in series with the motor to form a divider of sorts. Both of these are technically solutions, but have limitations imposed by how much heat they would inevitably produce. The safe operating range to prevent these systems from melting into their respective smoldering puddles of goo would be disappointingly small.

The Solution: The Humble Switch

Continuing with our heroic engineer performing his yard work in as complicated a fashion as possible, he has a mechanical load, that can essentially be described as a relatively large inertia and some friction, all driven by a DC motor, which is supplied by a battery. He would like to regulate the speed/torque of the motor so that his daughter is not thrown from the vehicle, but has run out of the usual tools he reaches for.

Perhaps, though, connecting the motor directly to the battery isn't such a bad idea after all… Especially if it is through a switch! The switch with the big red button in the salvaged parts bin would be perfect. Simply press the button until the desired speed is reached, then let off, and it will slow down. Repeat until the gravel reaches its destination. The large inertia of the car, the trailer, the gravel, and the tiny human will all make any changes in speed gradual as the torque is applied and removed. His daughter would only have to press and hold the button once every other second at the most to regulate the speed. Done! Gravel moved, work avoided.

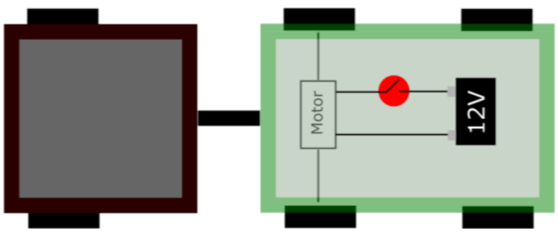

Figure 1: A Technical Schematic of a Kiddy Car

This method of pulsing a switch is the essence of power electronics, but usually the switches are a semiconductor such as a MOSFET and the filtering device is an inductor or capacitor rather than the load's inertia.

Pulse-Width Modulation

The ratio of "on" time to switching period, known as the duty ratio, dictates the steady-state average of the output. In the case of the engineer's fearless daughter enabling her father to goof off while claiming to work, the output is the voltage, and by extension, the speed. If the top speed of the system is 10MPH, and she holds down the switch for one second then coasts for one second, repeating the cycle over and over, then the duty ratio would be 0.5 and her average speed should be around 5MPH. If she were particularly fearless, she could hold the switch down for a second and a half then coast for a half second, arriving at an average speed of 7.5MPH. Reversing the times of the sequence gives the opposite result: on for a half second, off for 1.5s reduces the speed to 2.5MPH. Adjusting the switch's duty ratio in this fashion directly adjusts the average speed and is a special case of a technique referred to as pulse-width modulation, or PWM.

Of course, the fact that our engineer and young test pilot can regulate the average speed is not the only measure of the duo's success, which would become irritatingly apparent with each lunging and jolting motion. The output peaks and valleys of a system using PWM, known as the ripple and in this case the highest and lowest speeds during each cycle, are determined by the switching frequency and the size of the energy storage element. If she could press the button twice as often, the ripple would be reduced by half. Doubling the amount of gravel would roughly have the same effect, in addition to halving the workload!

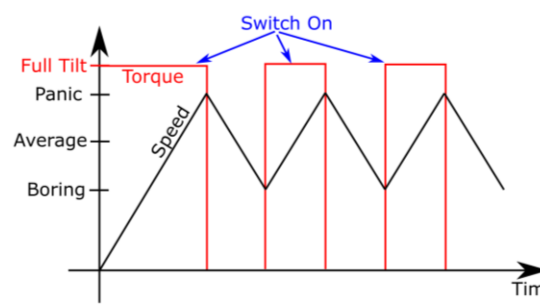

Figure 2: Primitive PWM

A Simple Application: The Buck Converter

Now let's say the engineer would like to add some colored, adjustable lighting to his daughter's room, but he has a bad case of do-it-yourself-itus. A circuit will need to be designed—naturally. LEDs are the choice of the day, but how to make them adjustable? Assuming we have a DC source squared away (refer to the application note on rectifiers), we again have a few choices available to us. A resistor is the crudest, most common way to regulate the current through an LED. It's reliable, simple, but inefficient. The engineer could even do something exotic like build a trans-impedance amplifier using an opamp, but again, inefficient. But perhaps… The switch! This time in the form of a MOSFET instead of a big red button. Connecting a timer circuit, perhaps a 555 timer, to the gate of a transistor allows the DC source to be connected directly to the string of LEDs, but only as often and for as long as the timer commands.

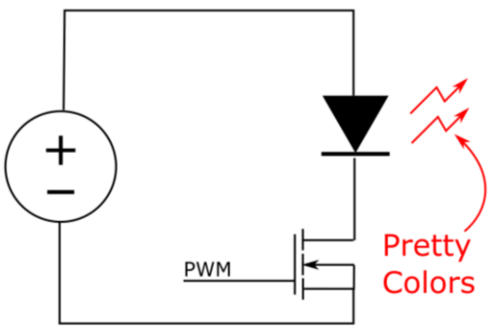

Figure 3: An Ideal Buck Converter (But a Bad Circuit Idea)

Soon enough our bumbling engineer will discover this is kind of a bad idea, but first let's talk about how it works in principle and address the problems once the circuit let's out its magic smoke. As with the janky kiddy vehicle, the *average* voltage will simply be the duty ratio times the DC source voltage, and the average current will be the product of the duty ratio and the operating current of the LEDs when connected to the source.

$$V_{ave} = DV_{DC}$$
$$I_{ave} = DI_{LED}$$

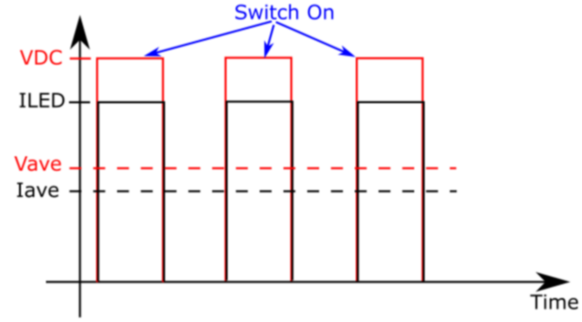

Figure 4: Ideal Buck Converter Waveforms

The average power to the load (the LEDs) in this system is then

$$P_{ave} = D^2 V_{DC} I_{LED}$$

This is pretty convenient—turn up the duty ratio and the power goes up which should translate directly to brightness. However, just like with the kiddy car, we are not only concerned with the average—however handy that might be to think about—but we are also concerned with the transient behavior. The first problem is that the switching frequency matters for comfortable perception of the light. Too slow and the gentle mood lighting our engineer was going for will end up more like the aircraft warning lights on a radio tower. Speeding up the frequency is clearly the answer, but if he doesn't increase it enough, his daughter could end up having the eerie feeling she is living in a stop-motion picture. At some frequency—on the order of 100Hz—we humans stop noticing the blinking, and instead see a continuous stream of light, the brightness of which is directly and conveniently related to the duty ratio.

But that is only the first hurdle. There is a chance that our somewhat careless engineer used a DC source with too high of a voltage—easy to do given the exponential relationship between LED current and voltage. Even though the *average* voltage and average current are functions of the duty ratio and can be held at seemingly reasonable magnitudes, the instantaneous voltage and current are not. They are simply the result of directly connecting the source to the LEDs. Poof! To solve this problem, we have a few options, which often appear in power electronics. The easy solution is to put a small resistor in series with the LEDs to limit the current. For small loads, this would be reasonable and is common for microcontroller connected LEDs. But… efficiency. For any sizeable amount of power flow this option is not viable because of the heat generated, or if we can find a way to jettison the heat, we would be flirting with being irresponsible. Besides, our engineer would at least like to have the option to power the prototype searchlight in his garage, and for this, he needs power. The other option, which is preferred in power electronics, is the passive filter. In this case, an inductor should perform this function just fine. There are some subtleties in switching strategy and inductor choice, but for now it is sufficient to say we can use the inductor to limit the ramp of the current in the same way the mass of the gravel limited the acceleration in the previous example. They are in fact mathematically identical.

As is typical in engineering though, the solution to one problem has inevitably introduced yet another problem. In the case of pulling the gravel, the vehicle was allowed to coast when the switch was opened, but when the MOSFET is opened after building up a current in the inductor, where will this current go? Opening the MOSFET would be analogous to calling a brick wall into existence in front of the kiddy car. In the car this would produce a spike in the force on the front, followed by broken plastic, scraped knees, and many tears. In the LED driver this would cause a spike in the voltage across the MOSFET, followed by a gentle pop, curling smoke, and again, many tears. For our engineer to make this work, he needs some way to make the circuit coast. Which brings us to the free-wheeling diode.

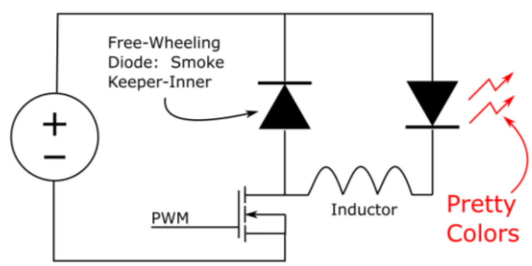

Figure 5: A More Reasonable Buck Converter

The free-wheeling diode in this case allows the current to continue flowing through the load without the source being connected, while ramping down the current, thus avoiding the sudden stop that would obliterate the innocent and unsuspecting MOSFET. This is a general problem that arises in power electronics as energy storage elements—inductors, capacitors, masses, springs, ect—must transfer energy in or out of their stores if we want them to change state. In the words of my old professor, if an engineer disobeys Kirchoff, he will rise up from the grave and strike you down. Dramatic, but true. However, what looks like an annoying bug actually turns out to be a feature. Where before we had one switch, we now have two and the result is more control and definition of the system's states which will turn out to be rather handy as the systems and PWM schemes become more complex and sophisticated.

Summary

The basic recipe for a power electronic system is switches (diodes, MOSFETs, IGBTs…) and energy storage devices (capacitors and inductors). The switches connect sources to loads and the energy storage act as filters that smooth the discrete states of switching out. The timing of the switches is controlled using a technique called pulse-width modulation to produce an output component we want—the average—but also to mitigate a component we don't want—the ripple.